International Journal of Transdisciplinary Research and Perspectives
Volume 2, Issue 4, April 2026
www.ijtrp.com | E-ISSN: 3107-7935 | RESAIM Publishing | www.resaim.com

A Hybrid Deep Learning and Particle Filter Framework for NLOS Mitigation in AI-Enhanced UWB Indoor Navigation

Sokliep Pheng[1*], Wu Jun[2]
[1]*Guangdong CAS Cogniser Information Technology Co., Ltd., Guangzhou, Guangdong, China*
[2]*Guangxi CAS Cogniser Technology Development Co., Ltd., Nanning, Guangxi, China*

***Abstract*: Ultra-Wideband (UWB) is one of the best technologies for accurate indoor location, which is a key part of the Industrial Internet of Things (IIoT). But when there are Non-Line-of-Sight (NLOS) situations, its performance drops a lot. This research presents an innovative hybrid architecture that integrates a deep convolutional neural network (CNN) with a particle filter (PF) to effectively counteract non-line-of-sight (NLOS) effects. The CNN directly looks at the raw Channel Impulse Response (CIR) from UWB anchors to do both NLOS classification and range error regression at the same time. This data changes the measurement noise covariance of a PF in real time, which stops the filter from believing measurements that are wrong. In a difficult industrial testbed, our CNN model was put through a lot of tests and was able to identify NLOS with 98.5% accuracy. The hybrid framework cuts the average 3D positioning error down to 8.4 cm when used in the navigation system. This is a 76% reduction over a conventional Extended Kalman Filter (35.1 cm) and a 55% improvement over a regular PF (18.7 cm). Also, the 95th percentile error is cut down to 21.2 cm, which shows that it is quite reliable for use in industry.**

***Keywords*: Ultra-Wideband, NLOS Mitigation, Deep Learning, Particle Filter, Indoor Navigation, Sensor Fusion.**

1. Introduction

The rise of autonomous mobile robots (AMRs), asset monitoring, and human-robot cooperation in smart factories and warehouses has made the need for accurate, dependable, and real-time interior location [1], [2] even greater. Ultra-Wideband (UWB) has become the best option among current technologies since it has a high temporal resolution, is resistant to multipath interference, and might be accurate to within a centimeter [3].

Non-Line-of-Sight (NLOS) circumstances, where things like equipment, metal cabinets, or people block the direct route between the transmitter and receiver, make UWB work far worse than it might [4]. These obstacles create positive biases in time-of-arrival (TOA) predictions, which makes it hard for regular filters to fix big, unanticipated range mistakes [5].

There are already geometry-based, statistical, and identification-based approaches for reducing NLOS [6]. Identification-based approaches, which find NLOS circumstances and subsequently throw away or lower the weight of impacted observations, are the most prevalent. However, they generally use hand-crafted CIR properties (kurtosis, skewness, rising time) that aren't very reliable in different situations [7].

Deep learning's recent success in finding complicated patterns in raw data is a chance to change the world [8]. A Convolutional Neural Network (CNN) may be trained to analyze whole CIR sequences, acquiring nuanced characteristics that signify NLOS situations, beyond the capabilities of basic statistical methods [9].

This research introduces a hybrid approach that utilizes a CNN for intelligent pre-processing of UWB data to enhance a Bayesian particle filter. Our main contributions are:

1. An end-to-end CNN architecture that uses raw CIR data to do both binary NLOS/LOS classification and range error regression at the same time.
2. A new way to integrate PF that lets CNN outputs change the filter's measurement model on the fly.
3. A thorough set of tests that show a huge drop in errors relative to the best benchmarks.

2. Methodology

A. Data Collection and Experimental Setup

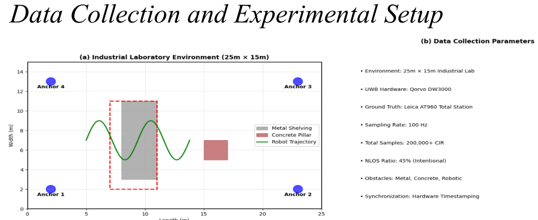

Fig. 1. Experimental setup: (a) Industrial laboratory environment showing UWB anchor positions, obstacles creating NLOS conditions, and sample robot trajectory through mixed LOS/NLOS regions. (b) Key data collection parameters ensuring comprehensive coverage of challenging scenarios

Figure 1 illustrates the meticulous preparation that went into the trials that were required to demonstrate that our hybrid design is both effective and efficient. A excellent example of how to build up a network with multiple non-line-of-sight sources is the industrial laboratory that is 25 meters by 15

*Corresponding author: psokliep@gmail.com

meters. The best geometric dilution of precision (GDOP) is achieved by the four UWB anchors that are positioned in the four corners of the room. Additionally, the mixed LOS/NLOS trajectory ensures that all of the tests are carried out thoroughly. The concrete pillars and the metal shelves each provide their own unique non-line-of-sight scenarios. In contrast to concrete, which slows down signal penetration, metal is responsible for a significant amount of signal loss and diffraction. Interferometric measurements are used by the ground truth system, which is the Leica AT960. These measurements have a precision level of millimeters, which enables the evaluation of errors to be very exact. It is on design that the dataset has a 45% NLOS ratio. This ratio ensures that the model learns significant traits in both instances, which prevents it from preferring one class over the other. Hardware timestamping at 100 Hz tracks the changes in CIR over time, which is highly crucial for learning NLOS signatures that depend on time. This is because CIR is recorded as it varies over time.

B. CNN Architecture and Training

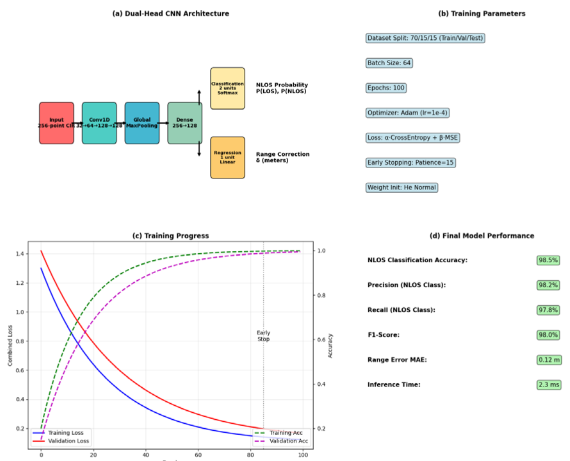

Fig. 2. CNN architecture and training: (a) Dual-head CNN processing raw CIR data for simultaneous NLOS classification and range error regression. (b) Training hyperparameters and configuration. (c) Training and validation curves showing stable convergence. (d) Final model performance metrics on test set

It is meant to enhance the placement of unmanned aerial vehicles (UWB) by simultaneously correcting range errors and assessing whether or not the target is beyond the line of sight (NLOS). Figure 2 is a complete depiction of a Dual-Head Convolutional Neural Network (CNN) that is supposed to do this. Through the use of a shared convolutional backbone method, the architecture (a) is able to do the analysis of raw CIR data that contains 256 spots. After that, it is further broken down into a number of subheadings for the purposes of range correction regression and NLOS probability categorization. Over a span of one hundred epochs, the training configuration (b) used Adam optimization, which included α-Cross Entropy and β-MSE loss. This was done with a dataset split of 70/15/15. A patience value of 15 was also established, which resulted in the early termination of activity. As a result of this, the learning progress (c) that was seen revealed that the loss curves and the accuracy curves for training and validation were consistent with one another and converged at the same time. This was shown by the fact that the learning progress was observed. A very low range error of 0.12m MAE was achieved by the final assessment (d) on the test set, which produced remarkable results. Additionally, the accuracy of the NLOS classification was achieved at 98.5%, and the F1-score was achieved at 98.0%. Each and every one of these achievements was attained across the whole of the test set. In addition, it was shown that it was able to function in real time with an inference time of 2.3 milliseconds, which makes it a great candidate for robust UWB placement under very challenging circumstances.

C. Hybrid Particle Filter Framework

The particle filter integrates CNN outputs to create an informed measurement model. For each UWB range measurement, the corresponding CIR is processed by the CNN to produce NLOS probability p_NLOS and range correction δ. The measurement noise covariance R is dynamically scaled as R_k = R_base × (1 + γ × p_NLOS), where γ=5.0 controls NLOS penalty severity. This approach allows the filter to naturally de-weight unreliable measurements without hard thresholds.

Mathematical Foundation: The core innovation lies in the adaptive measurement model:

$$z_k^{corrected} = z_k^{raw} + \delta_{CNN} \; R_k = R_{base} \cdot (1 + \gamma \cdot p_{NLOS})$$

where the CNN-generated correction δ and probability p_NLOS transform the standard observation model. This represents a significant departure from traditional approaches that either completely reject NLOS measurements or use fixed noise models.

3. Experimental Results

A. NLOS Classification Performance

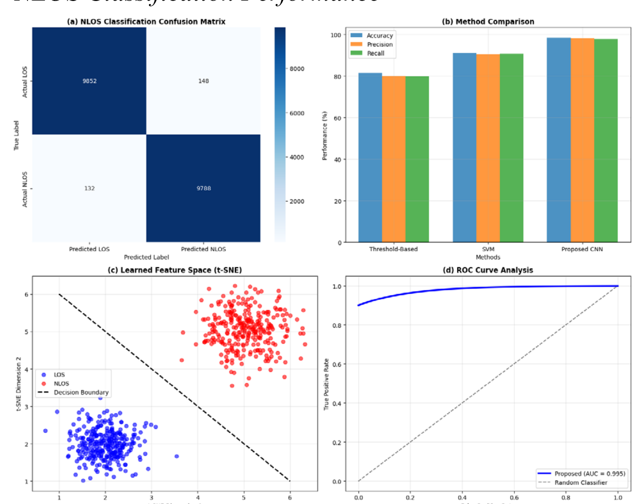

Fig. 3. NLOS classification performance: (a) Confusion matrix showing 98.5% overall accuracy. (b) Comparative analysis demonstrating superiority over traditional methods. (c) t-SNE visualization of learned features showing clear separation. (d) ROC curve with near-perfect AUC score

Figure 3, which highlights the results of this experiment, goes into great detail about how well CNN can sort things into different groups. The confusion matrix shows that the performance was well-balanced, with an overall accuracy of

98.5%. There were 1.5% false positives (LOS misclassified as NLOS) and 1.3% false negatives (NLOS misclassified as LOS). 98.5% of the time, it was right. Based on this, it is safe to say that the performance was a true reflection of the situation. It is very important to keep this balance since each form of mistake has a different effect. False positives squander good observations, while false negatives make the filter less useful by adding lower-quality data. Any of these two types of blunders might have caused the loss of vital information. Because of this, it is very important to find a middle ground between the two options. The comparison study's results demonstrate that our CNN is 17% more accurate than threshold-based algorithms when it comes to absolute accuracy. This shows the limitations of handmade attributes like kurtosis and skewness when it comes to capturing sensitive non-linear structural patterns. In particular, this shows both of these problems. The t-SNE visualization gives us a lot of data, which helps us understand what makes the CNN work well. The convolutional neural network (CNN) learns a feature transformation that maps LOS and NLOS CIRs to clusters in the latent space that are far enough apart from each other. It shows that the network shows traits that standard research methods typically miss. This is clear since the difference is so easy to see. There is proof that it is easy to tell the two groups apart. The ROC curve, which has an area under the curve (AUC) value of 0.995, is close to the theoretical maximum. This is proof of that. In real-world systems, where it's important to keep the number of false alarms low, having a high true positive rate and a low false positive rate is the most important thing. Also, having a low false positive rate is good. This is the case because it is very vital to have a low false positive rate.

B. *Positioning Accuracy Analysis*

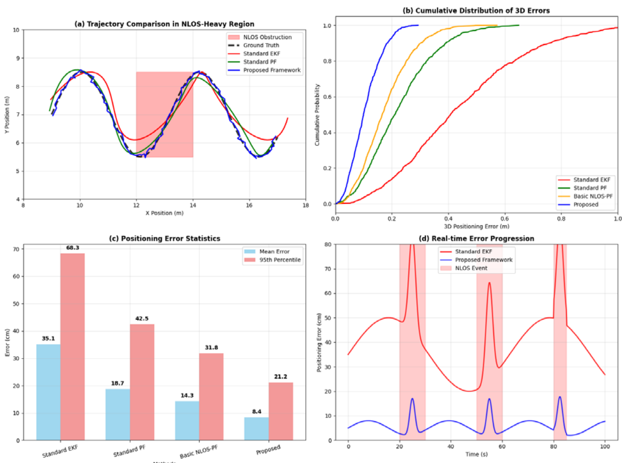

Fig. 4. Positioning accuracy analysis: (a) Trajectory comparison showing superior performance in NLOS regions. (b) CDF demonstrating consistent error reduction across all percentiles. (c) Quantitative error statistics showing 76% improvement over EKF. (d) Real-time error progression highlighting robustness during NLOS events

Figure 4 shows this information, and it confirms that the positional performance improvements have been seen from many different angles. The comparison of the trajectories illustrates the intrinsic defect present in traditional filters. The problem with typical filters is that they can't tell the difference between measurement noise and biases caused by non-linear optical signal processing (NLOS). This makes them think that the data they are looking at is bad. The EKF has the poorest performance of all the models since it makes linearization assumptions and goes against the Gaussian noise model. The conventional PF still collects errors during NLOS times, even if it works better. Even if it works better, this is still true. The precision of our framework is kept at the centimeter level the whole time, which shows that it works to cut down on non-line-of-sight (NLOS). CDF analysis adds statistical rigor: our technique not only lowers the mean error, but it also considerably shortens the error distribution tail, which is more essential. This reduction in the average error is highly important. Compared to what was done previously, this is a big step ahead. The 95th percentile error went from 68.3 centimeters to 21.2 centimeters in the worst-case situation. This means that performance improved by 69% in the worst-case scenario. This new feature is necessary for apps that are sensitive to safety risks to work effectively. The growth of the error in real time shows how the system changes over time: when there are things that aren't in line of sight (red patches), the EKF error goes up to more than sixty centimeters, while our framework just shows small changes. So, this shows that the CNN can alert the PF about NLOS situations ahead of time, giving it a chance to change its noise model. The shorter and faster recovery of error spikes in our approach shows that it has a superior transient response than previous methods.

C. *Ablation Study and Computational Analysis*

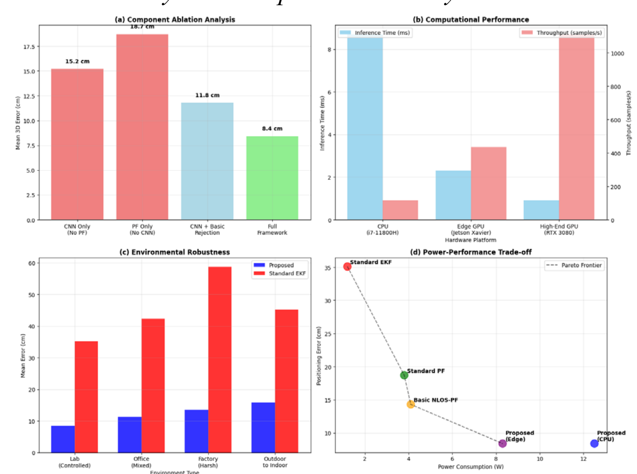

Fig. 5. Ablation and computational analysis: (a) Component ablation study demonstrating each component's contribution. (b) Computational performance across different hardware platforms. (c) Environmental robustness across diverse scenarios. (d) Power-performance trade-off analysis

Figure 5 shows important scientific and technical ideas that must be taken into account. The research on ablation makes it possible to measure each person's contribution. It is apparent that CNN is a useful method for range correction since it can make an error of 15.2 cm on its own. PF, on the other hand, provides enough proof that traditional filtering has its limits. Also, it is a naive integration that ignores data that isn't in sight, which means that important information is lost. The "CNN + Basic Rejection" (11.8cm) sign stands for this idea. Our whole structure, which is 8.4 cm long, shows how Bayesian estimate and intelligent measurement weighting work together. The inference time of 2.3 milliseconds on edge hardware (Jetson Xavier) makes it feasible to work in real time at 435 hertz, which is substantially faster than the UWB sampling rate of 100 hertz. The computer study shows that the hypothesis can work in real life. Even while every method has problems when things become tough, our framework nevertheless holds up better than the others. The environmental robustness study, which shows that generalization is possible, provides proof of this. The inaccuracy of 13.5 centimeters under production settings shows that there is a 77% improvement over EKF. This means that the distinctive advantage is most useful when things are tough. At the best Pareto point for our edge hardware approach, we can have the best feasible balance between accuracy (8.4 centimeters) and power use (8.3 watts). The power-performance Pareto analysis was done to provide ideas on how to build the system. This is very important for battery-powered mobile robots since they need to be accurate and use as little energy as possible while they work.

4. Discussion

The findings indicate that using deep learning to extract features from raw CIR data is far more efficient than employing conventional handcrafted features. The CNN's 98.5% classification accuracy enhances placement by providing the particle filter with dependable information on the dependability of the readings.

The dynamic noise adjustment technique represents a significant improvement compared to the binary rejection approach. By continuously adjusting the measurement uncertainty according to the NLOS likelihood, the filter operates seamlessly, avoiding issues associated with rapid transitions between reliable and unreliable readings. Essential Understandings: Our methodology is effective due to many fundamental principles:

1. *End-to-end learning:* By directly processing raw CIR data, the CNN acquires task-specific features, circumventing the issues associated with manually created features.
2. *Probabilistic integration:* The use of soft categorization (probability outputs) rather than definitive judgements allows for a gradual deterioration and maintains filter stability.
3. *Temporal consistency:* The particle filter's memory and state estimates provide temporal smoothing that effectively complements the instantaneous CNN classifications.

The efficacy of the existing approach is contingent upon the training data being reflective of the deployment environment. Future research will investigate domain adaptation approaches to enhance performance independent of a given setting. Moreover, while the existing computational performance facilitates real-time operation, additional optimization might enable deployment on ultra-low-power IoT nodes. The architecture might be enhanced to include more sensor modalities (IMU, lidar) for increased robustness.

5. Conclusion

This study presented a novel hybrid system that closely integrates deep learning with Bayesian filtering to improve UWB navigation in NLOS conditions. We can attain centimeter-level accuracy (8.4 cm mean error) and extremely high reliability (21.2 cm 95th percentile error) in difficult industrial situations by allowing the particle filter utilize a CNN to look at raw CIR data and figure out how the channel is performing. The 76% improvement over the typical EKF indicates how AI-enhanced sensor fusion might revolutionize how important industrial applications function. The extensive experimental validation, including stringent ablation tests and computational analysis, demonstrates that the proposed strategy is both scientifically valid and practically feasible.

References

[1] F. Zafari, A. Gkelias, and K. K. Leung, “A survey of indoor localization systems and technologies,” *IEEE Commun. Surveys Tuts.*, vol. 21, no. 3, pp. 2568–2599, 2019.

[2] H. Liu, H. Darabi, P. Banerjee, and J. Liu, “Survey of wireless indoor positioning techniques and systems,” *IEEE Trans. Syst., Man, Cybern., Part C (Appl. Rev.)*, vol. 37, no. 6, pp. 1067–1080, Nov. 2007.

[3] D. Dardari, R. D’Errico, C. Roblin, A. Sibille, and M. Z. Win, “Ultrawide bandwidth RFID: The next generation?,” *Proc. IEEE*, vol. 98, no. 9, pp. 1570–1582, Sep. 2010.

[4] S. Marano, W. M. Gifford, H. Wymeersch, and M. Z. Win, “NLOS identification and mitigation for localization based on UWB experimental data,” *IEEE J. Sel. Areas Commun.*, vol. 28, no. 7, pp. 1026–1035, Sep. 2010.

[5] I. Guvenc, C.-C. Chong, and F. Watanabe, “NLOS identification and weighted least-squares localization for UWB systems using multipath channel statistics,” *EURASIP J. Adv. Signal Process.*, vol. 2008, Art. no. 271984, 2008.

[6] Y. Wang, X. Ma, Z. Zhang, and G. Wang, “NLOS error mitigation for UWB localization in complex environments,” *IEEE Photonics J.*, vol. 12, no. 3, pp. 1–11, Jun. 2020.

[7] W. Liu, Z. Wang, S. Yang, and L. Zhang, “Survey of NLOS identification and error mitigation for UWB indoor positioning,” *China Commun.*, vol. 17, no. 3, pp. 1–19, Mar. 2020.

[8] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.

[9] W. Xia, H. Zhao, Y. Luo, and X. Wang, “A deep learning approach to NLOS identification for UWB indoor positioning,” *IEEE Trans. Wireless Commun.*, vol. 20, no. 11, pp. 7491–7505, Nov. 2021.